

# SRI International



## MODELING AND USING PHYSICAL CONSTRAINTS IN SCENE ANALYSIS

Technical Note 267

September 1982

By: Martin A. Fischler, Stephen T. Barnard,  
Robert C. Bolles, Michael Lowry, Lynn Quam,\*  
Grahame Smith, Andrew Witkin\*

Artificial Intelligence Center  
Computer Science and Technology Division

The research reported herein was supported by the Defense Advanced Research Projects Agency under Contract Nos. MDA903-79-C-0588 and DAAG29-79-C-0216; these contracts are monitored by the U. S. Army Engineer Topographic Laboratory and by the U. S. Army Research Office.

---

\* Current affiliation: Fairchild Artificial Intelligence Laboratory, Palo Alto, California.

## ABSTRACT

This paper describes the results obtained in a research program ultimately concerned with deriving a physical sketch of a scene from one or more images. Our approach involves modeling physically meaningful information that can be used to constrain the interpretation process, as well as modeling the actual scene content. In particular, we address the problems of modeling the imaging process (camera and illumination), the scene geometry (edge classification and surface reconstruction), and elements of scene content (material composition and skyline delineation).

## I INTRODUCTION

Images are inherently ambiguous representations of the scenes they depict: images are 2-D views of 3-D space, they are single slices in time of ongoing physical and semantic processes, and the light waves from which the images are constructed convey limited information about the surfaces from which these waves are reflected. Therefore, interpretation cannot be strictly based on information contained in the image; it must involve, additionally, some combination of a priori models, constraints, and assumptions. In current machine-vision systems this additional information is usually not made explicit as part of the machine's data base, but rather resides in the human operator who chooses the particular techniques and parameter settings to reflect his understanding of the scene context. This paper describes a portion of the SRI program in machine vision research that is concerned with identifying and modeling physically meaningful information that can be used to automatically constrain the interpretation process. In particular, as an adjunct to any autonomous system with a generalized competence to analyze imaged data of 3-D real-world scenes, we believe that it is necessary to explicitly model and use the following types of knowledge:

- (1) Camera model and geometric constraints (location and orientation in space from which the image was acquired, vanishing points, ground plane, geometric horizon, geometric distortion).
- (2) Photometric and illumination models (atmospheric and image-processing system intensity-transfer functions, location and spectrum of sources of illumination, shadows, highlights).
- (3) Physical surface models (description of the 3-D geometry and physical characteristics of the visible surfaces; e.g., orientation, depth, reflectance, material composition).
- (4) Edge classification (physical nature of detected edges; e.g., occlusion edge, shadow edge, surface intersection edge, material boundary edge, surface marking edge).
- (5) Delineation of the visible horizon (skyline)
- (6) Semantic context (e.g., urban or rural scene, presence of roads, buildings, forests, mountains, clouds, large water bodies, etc.).

In the remainder of this paper, we will describe in greater detail the nature of the above models, our research results concerning how the parameters for some of these models can be automatically derived from image data, and how the models can be used to constrain the interpretation process in such tasks as stereo compilation and image matching.

If we categorize constraints according to the scope of their influence, then the work we describe is primarily concerned with global and extended constraints rather than with constraints having only a local influence. To the extent that constraints can be categorized as geometric, photometric, or semantic and scene dependent, it would appear that we have made the most progress in understanding and modeling the geometric constraints.

## II CAMERA MODELS AND GEOMETRIC CONSTRAINTS

The camera model describes the relationship between the imaging device and the scene; e.g., where the camera is in the scene, where it is looking, and more specifically, the precise mapping from points in the scene to points in the image. In attempting to match two views of the same scene taken from different locations in space, the camera model provides essential information needed to contend with the projective differences between the resulting images.

In the case of stereo reconstruction, where depth (the distance from the camera to a point in the scene) is determined by finding the corresponding scene point in the two images and using triangulation, the camera models (or more precisely, the relative camera model) limit the search for corresponding points to one dimension in the image via the "epipolar" constraint. The plane passing through a given scene point and the two lens centers intersects the two image planes along straight lines; thus a point in one image must lie along the corresponding (epipolar) line in the second image, and one need only search along this line, rather than the whole image to find a match.

When human interaction is permissible, the camera model can be found by having the human identify a number of corresponding points in the two images and using a least-squares technique to solve for the parameters of the model [5]. If finding the corresponding points must be carried out without human intervention, then the differences in appearance of local features from the two viewpoints will cause a significant percentage of false matches to be made; under these conditions, least squares is not a reliable method for model fitting. Our approach to this problem [3] is based on a philosophy directly opposite to that of least-squares -- rather than using the full collection of matches in an attempt to "average out" errors in the model-fitting process, we randomly select the smallest number of points needed to solve for the camera model and then enlarge this set with additional correspondences that are compatible with the derived model.

If the size of the enlarged compatibility set is greater than a bound determined by simple statistical arguments, the resulting point set is passed to a least-squares routine for a more precise solution. We have been able to show that as few as three correspondences are sufficient to directly solve for the camera parameters when the three-space relationships of the corresponding points are known; a recent result [13] indicates that 5 to 8 points are necessary to solve for the relative camera model parameters when three space information is not available a priori.

The perspective imaging process (the formation of images by lenses) introduces global constraints that are independent of the explicit availability of a camera model; particularly important are the detection and use of "vanishing points." A set of parallel lines in 3-D space, such as the vertical edges of buildings in an urban scene, will project onto the image plane as a set of straight lines intersecting at a common point. Thus, for example, if we can locate the vertical vanishing point, we can strongly constrain the search for vertical objects such as telephone or power poles or building edges, and we can also verify conjectures about the 3-D geometric configuration of objects with straight edges by observing which vanishing points these edges pass through. The two horizontal vanishing points corresponding to the rectangular layout of urban areas, the vanishing point associated with a point of illumination [8], and the vanishing point of shadow edges projected onto a plane surface in the scene, provide additional constraints with special semantic significance. The detection of clusters of straight parallel lines by finding their vanishing points can also be used to automatically screen large amounts of imagery for the presence of man-made structures.

The technique we have employed to detect potential vanishing points involves local edge detection by finding zero-crossings in the image convolved with both Gaussian and Laplacian operators [9], fitting straight line segments to the closed zero-crossing contours, and then finding clusters of intersection points of these straight lines. In

order to avoid the combinatorial problem of computing intersection points for all pairs of lines, or the even more unreasonable approach of plotting the infinite extension of all detected line segments and noting those locations where they cluster, we have implemented the following technique. Consider a unit radius sphere physically positioned in space somewhere over the image plane (there are certain advantages to locating the center of the sphere at the camera focal point if this is known, in which case it becomes the Gaussian sphere [6,7], but any location is acceptable for the purpose under consideration here). Each line segment in the image plane and the center of the sphere define a plane that intersects the sphere in a great circle -- if two or more straight lines intersect at the same point on the image plane, their great circles will intersect at two common points on the surface of the sphere, and the line passing through the center of the sphere and the two intersection points on the surface of the sphere will also pass through the intersection point in the image plane.

### III EDGE CLASSIFICATION

An intensity discontinuity in an image can correspond to many different physical events in the scene, some very significant for a particular purpose, and some merely confusing artifacts. For example, in matching two images taken under different lighting conditions, we would not want to use shadow edges as features; on the other hand, shadow edges are very important cues in looking for (say) thin raised objects. In stereo matching, occlusion edges are boundaries that area correlation patches should not cross (there will also be a region on the "far" side of an occlusion edge in which no matches can be found); occlusion edges also define a natural distance progression in an image even in the absence of stereo information. If it is possible to assign labels to detected edges describing their physical nature, then those interpretation processes that use them can be made much more robust.

We have implemented an approach to detecting and identifying both shadow and occlusion edges, based on the following general assumptions about images of real scenes:

- (1) The major portion of the area in an image (at some reasonable resolution for interpretation) represents continuous surfaces.
- (2) Spatially separated parts of a scene are independent, and their image projections are therefore uncorrelated.
- (3) Nature does not conspire to fool us; if some systematic effect is observed that we normally would anticipate as caused by an expected phenomena due to imaging or lighting, then it is likely that our expectations provide the correct explanation; e.g., coherence in the image reflects real coherence in the scene, rather than a coincidence of the structure and alignment of distinct scene constituents.

Consider a curve overlayed on an image as representing the location of a potential occlusion edge in the scene. If we construct a series of curves parallel to the given one, then we would expect that for an occlusion edge, there would be a high correlation between adjacent curves on both sides of the given curve, but not across this curve. That is, on each side, the surface continuity assumption should produce the required correlation, but across the reference curve the assumption of remote parts of the scene being independent should produce a low correlation score. In a case where the reference curve overlays a shadow edge, we would expect a continuous high (normalized) correlation between adjacent curves on both sides and across the reference curve, but the regression coefficients should show a discontinuity as we cross the reference curve. This technique is described in greater detail in [14]. Figures 1 and 2 show experimental results for shadow and occlusion edges.

#### IV INTENSITY MODELING (and Material Classification)

Given that there is a reasonably consistent transform between surface reflectance and image intensity, the exact nature of this transform is not required to recover rather extensive information about the geometric configuration of the scene. It is even reasonable to assume that shadows and highlights can be detected without more precise knowledge of the intensity mapping from surface to image; but if we wish to recover information about actual surface reflectance or physical composition of the scene, then the problem of intensity modeling must be addressed.

Even relatively simple intensity modeling must address three issues: (1) the relationship between the incident and reflected light from the surface of an object in the scene as a function of the material composition and orientation of the surface; (2) the light that reaches the camera lens from sources other than the surface being viewed (e.g., light reflected from the atmosphere); and (3) the relationship between the light reaching the film surface and the intensity value ultimately recorded in the digital image array.

Our approach to intensity modeling assumes that we have no scene-specific information available to us other than the image data. We use a model of the imaging process that incorporates our knowledge of the behavior of the recording medium, the properties of atmospheric transmission, and the reflective properties of the scene materials. For aerial imagery we use an atmospheric model that assumes a constant amount of light, (independent of scene radiance), is scattered by the atmosphere into the camera.

$$I=R+S$$

where  $I$  is the image irradiance,  $R$  the scene radiance, and  $S$  the image irradiance caused by atmospheric scattering. We use a logarithmic relationship between image irradiance and film density  $D$ ,



$$D=a*\log(I) + d$$

where  $a$  and  $d$  are film constants, whose values need to be calculated. For a surface radiance model we assume Lambertian behaviour (the reflected light is proportional to the incident light, the constant of proportionality is a function of the surface material, and the relative brightness of the surface is independent of the location of the viewer).

$$R=EAN$$

where  $E$  is the illumination strength (scene irradiance),  $A$  the surface reflection or albedo, and  $N$  a function related to the effects of surface orientation (for Lambertian surfaces this is a function of the angle between the surface normal and the light direction).

If for the present we ignore surface orientation effects, that is we assume all surfaces are orientated in the same direction, then our model has the form

$$D=a*\log(A+b)+c$$

where  $a, b$ , and  $c$  are constants that need to be determined.  $b$  is the ratio of atmospheric scattering to illumination irradiance.

We calibrate our model by identifying a few regions of known material in an image. Three materials are sufficient. The fitting is achieved by guessing  $b$  - we know  $b$  lies in the range 0 to 1 - applying the least squares method to the resultant linear equation to calculate  $a, c$ , and the residual sum, and adjusting  $b$  to minimize this residual sum.

The resultant model is used to transform the given image into a new image depicting the scene albedo. The albedo image has been used to provide an initial classification (and partitioning) of the scene using straight forward classification techniques based on "known" surface albedos. This technique allows classification without the need to provide training samples of all classes that are present in the image.

## V SHADOW DETECTION (and Raised Object Cueing)

The ability to detect and properly identify shadows is a major asset in scene analysis. For certain types of features, such as thin raised objects in a vertical aerial image, it is often the case that only the shadow is visible. Knowledge of the sun's location and shadow dimensions frequently allows us to recover geometric information about the 3-D structure of the objects casting the shadows, even in the absence of stereo data [8,10]; but perhaps just as important, distinguishing shadows from other intensity variations eliminates a major source of confusion in the interpretation process.

Given an intensity discontinuity in an image, we can employ the edge labeling technique described earlier to determine if it is a shadow edge. However, some thin shadow edges are difficult to find, and if there are lots of edges, we might not want to have to test all of them to locate the shadows. We have developed a number of techniques for locating shadow edges directly, and will now describe a simple but effective method for finding the shadows cast by thin raised objects (and thus locating the objects as well).

We assume we either know the approximate sun direction, or equivalently, the shadow vanishing point. We first employ a thin line detector oriented parallel to the sun direction at every location in the image, and then apply a moving-window averaging technique in the sun's direction to further enhance the line detector's response and reduce noise. The result of these operations is to smear both the noise and the thin shadow lines. We next thin the shadow lines, eliminate all weak responses, and overlay the result on the original image. The foot of each shadow line now points to the base of the thin raised object casting the shadow. Given the results from two (or more) images taken at different times, the intersections of shadow lines locates the objects more precisely and also eliminates false alarms.

The same technique has been applied to the detection of raised objects of extended size. Shadow edges of the extended object are detected and used to locate the object. Figures 3-10 show this approach to detecting both thin and extended raised objects.

## VI VISUAL SKYLINE DELINEATION

Although not always a well defined problem, delineation of the land-sky boundary provides important constraining information for further analysis of the image. Its very existence in an image tells us something about the location of the camera relative to the scene (i.e., that the scene is being viewed at a high-oblique angle), allows us to estimate visibility (i.e., how far we can see -- both as a function of atmospheric viewing conditions, and as a function of the scene content), provides a source of good landmarks for (autonomous) navigation, and defines the boundary beyond which the image no longer depicts portions of the scene having fixed geometric structure.

In our analysis, we generally assume that we have a single right-side-up image in which a (remote) skyline is present. Confusing factors include clouds, haze, snow-covered land structures, close-in raised objects, and bright buildings or rocks that have intensity values identical to those of the sky (a casual inspection of an image will often provide a misleading opinion about the difficulty of skyline delineation for the given case). Our initial approach to this problem was to investigate the use of slightly modified methods for linear delineation [4] and histogram partitioning based on intensity and texture measures; we employ fairly simple models of the relationship between land, sky, and cloud brightness and texture.

Currently, we are employing a region based technique which operates as follows:

To eliminate spurious regions and gaps in region boundaries caused by noise we first reduce the given image by a factor of at least 4. We

partition the image into a nested pyramid of regions; each region being one in which every pixel has an intensity value which differs by less than some given threshold, from a least one of its 4-neighbors. The nested pyramid is constructed by using a sequence of increasing threshold values (e.g. 2,4,8,...); thus if  $T_1$  and  $T_2$  are thresholds such that  $T_1 < T_2$ , then any region found with threshold  $T_1$  is necessarily identical to a subregion or a region found with threshold  $T_2$ .

A "sky seed" is found by identifying the region that dominates the very top of the picture with a segmentation threshold of 2 (this is the lowest threshold that allows a gradient to exist within a region). For a clear sky, or a sky with cumulus clouds completely surrounded by clear sky, this step usually identifies the entire sky. Figure 11 shows an urban scene with overcast sky and figure 12 shows the same scene with the sky seed overlayed.

As an additional piece of information, the sky seed is classified as clear sky, overcast sky, or patchy clouds. Patchy cumulus clouds appear as large bright regions within a clear sky region, while the brightness function for a clear sky can be modeled as a linear function of the image coordinates. Although the equations governing clear sky luminance are complex integro-differential equations, it was determined empirically that for the viewing angles produced by a 50mm lens, a (linear) planar model provided a good fit. To determine whether the sky seed is clear sky or overcast sky, a least squares fit to the planar model is made, and the mean square error, corrected by the measured intensity variance, is compared to a fixed thresholded. The classification of the sky into clear/overcast/patchy clouds can help to resolve some of the confusing factors in skyline detection, but this information is not currently used.

Next, a line spanning the picture from right to left is found that is either at or below the true skyline; this line is found by doubling the threshold for segmentation until the region containing the sky seed touches the bottom of the picture. Since we make an initial assumption

that the sky does not touch the bottom of the picture, this threshold is then backed off by a factor of 2 and a "land seed" is defined as the complement of the sky region. The assumption here is that the skyline, or some boundary in the land region, is of higher contrast than any extended boundary within the sky. For all 15 test pictures that we employed in our experiments, this assumption was only violated once (where a particularly bright cumulus cloud on the horizon formed a brighter boundary with the sky than a bright rock on the horizon; such a situation can be easily detected after initial processing). Figure 13 shows the case in which a region containing the sky seed touches the bottom of the picture at a threshold of 16, and Figure 14 shows the picture split into a sky seed, land seed, and ambiguous unclassified portion. The land seed is determined by using a threshold of 8. Figure 15 shows an additional and more typical example of skyline delineation.

In a substantial number of pictures the sky and land seeds touch, thereby delineating the skyline. If the sky and land seeds do not have a common boundary, a portion of the picture is left unclassified, bounded by the sky seed above and the land region below. Current work focuses on developing methods to disambiguate the unclassified portion of the picture. The methods under development are generic to all types of scenes and our approach does not use semantic knowledge of particular land features. Prior work on this topic, employing considerable semantic knowledge, is contained in Sloan [11].

## VII SURFACE MODELING

Obtaining a detailed representation of the visible surfaces of the scene, as (say) a set of point arrays depicting surface orientation, depth, reflectance, material composition, etc., is possible from even a single black and white image [12,2]. A large body of work now exists on this topic,(see [15,16] for recent work by our group), and although directly relevant to our efforts, it is not practical to attempt a discussion of this material here. There is, however, one key difference

between surface modeling and the other topics we have discussed -- the extent to which the particular physical knowledge modeled constrains the analysis of other parts of the scene. In this paper we have been primarily concerned with physical models that provide global or extended constraints on the analysis; surface modeling via point arrays provides a very localized constraining influence.

## VIII CONSTRAINT-BASED STEREO COMPILATION

The computational stereo paradigm encompasses many of the important task domains currently being addressed by the machine-vision research community [1]; it is also the key to an application area of significant commercial and military importance -- automated stereo compilation. Conventional approaches to stereo compilation, based on finding dense matches in a stereo image pair by area correlation, fail to provide acceptable performance in the presence of the following conditions typically encountered in mapping cultural or urban sites: widely separated views (in space or time), wide angle views, oblique views, occlusions, featureless areas, repeated or periodic structures. As an integrative focus for our research, and because of its potential to deal with the factors that cause failure in the conventional approach, we are constructing a constraint-based stereo system that encompasses many of the physical modeling techniques discussed above.

Figure 16 show how a stereo system can exploit global geometric constraints. First, straight lines and vanishing points are found in the two stereo images as described earlier (see Section II). Lines are first classified according to which vanishing point they pass through. Those lines not associated with the detected vanishing points are ignored. The vanishing points in the two stereo views are then matched. The direction in space established by a vanishing point is a feature of the scene which is invariant under translation of the camera. Two matches of vanishing points are sufficient to calculate the rotational differences between the cameras i.e., the rotation required to bring one

camera's vanishing points into congruence with the other's. Two matches of ordinary points are now sufficient to determine the translation of one camera with respect to the other (up to an unknown scaling factor).

Using vanishing points can improve stereo matching even when the exact camera model is unknown. In Figures 16i and 16j, lines passing through a vanishing point in one image are first matched to the set of lines passing through the corresponding vanishing point in the other image. For example, right-image lines passing through the vertical vanishing point are only matched to left-image lines that also pass through the vertical vanishing point. Within these subsets, lines are matched according to a score based on four features: (1) difference in distance from the vanishing point to the lines, (2) ratio of lengths, (3) difference in contrast and, (4) difference in phase i.e., the angle the line makes with the image-horizontal. Each subscore is a value in the interval  $[0,1]$ . The value represents the likelihood of this combination of the four features. The subscores are combined multiplicatively, and the combination with the maximum score (above a preset threshold) is chosen. Even this simple matching technique, using no search or relaxation, finds an adequate number of correct matches.

## IX CONCLUDING COMMENTS

When a person views a scene, he has an appreciation of where he is relative to the scene, which way is up, the general geometric configuration of the surfaces (especially the support and barrier surfaces), and the overall semantic context of the scene. The research effort we have described is intended to provide similar information to constrain the more detailed interpretation requirements of machine vision (e.g., such tasks as stereo compilation and image matching).

## REFERENCES

1. S. T. Barnard and M. A. Fischler, "Computational Stereo," ACM Surveys, 1982 (in press).
2. H. G. Barrow and J. M. Tenenbaum, "Recovering intrinsic scene characteristics from images," in Computer Vision Systems (A. Hanson and E. Riseman ed.) Academic Press, pp. 3-26 (1978).
3. M. A. Fischler and R. C. Bolles, "Random Sample Consensus: A paradigm for model fitting with applications to image analysis and automated cartography," CACM, Vol. 24(6), pp. 381-395 (June 1981).
4. M. A. Fischler, J. M. Tenenbaum, and H. C. Wolf, "Detection of roads and linear structures in low-resolution aerial imagery using a multisource knowledge integration technique," Computer Graphics and Image Processing, Vol. 15(3), pp. 201-223 (March 1981).
5. D. B. Gennery, "Least-squares stereo-camera calibration," Stanford Artificial Intelligence Project Internal Memo, Stanford University (1975).
6. D. A. Huffman, "Realizable configuration of lines in pictures of polyhedra," in Machine Intelligence (Elcock and Michie, ed.), Edinburgh University Press, Edinburgh, Scotland, pp. 493-509 (1977).
7. J. R. Kender, "Shape from Texture," (Ph.D. Thesis, Report No. CMU-CS-81-102) Carnegie-Mellon University, Pittsburgh, Pennsylvania (November 1980).
8. D. G. Lowe and T. O. Binford, "The interpretation of three-dimensional structure from image curves," IJCAI-81, pp. 613-618 (1981).
9. D. Marr and E. C. Hildreth, "Theory of edge detection," MIT AI Lab Memo 518 (1979).
10. S. Shafer and T. Kanade, "Using shadows in finding surface orientations," (Report No. CMU-CS-82-100) Carnegie-Mellon University, Pittsburgh, Pennsylvania (January 1982).
11. K. R. Sloan, Jr., "World model driven recognition of natural scenes," University of Pennsylvania, Philadelphia, Pennsylvania (June 1977).
12. J. M. Tenenbaum, M. A. Fischler, and H. G. Barrow, "Scene Modeling: A structural basis for image description," Computer Graphics and Image Processing, Vol. 12(4), pp. 407-425 (April 1980).



13. R. Y. Tsai and T. S. Huang, "Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces," University of Illinois, Urbana, Illinois (August 1981).
14. A. Witkin, "Recovering intrinsic scene characteristics from images," SRI Project 1019, Interim Technical Report, SRI International, Menlo Park, California (September 1981).
15. A. P. Pentland, "Local computation of shape," Proceedings AAAI, (1982)
16. G. B. Smith, "The recovery of surface orientation from image irradiance," Image Understanding Workshop, Stanford (September 1982).





FIGURE 1 EXAMPLE OF CAST-SHADOW EDGE  
(Edge Classification)






FIGURE 2 EXAMPLE OF EXTREMAL EDGE  
 (Edge Classification)

● MASK



APPLIED AT RIGHT ANGLES  
TO SHADOW DIRECTION



● SCORE = MINIMUM (a,c) - b

● HIGH SCORE IMPLIES LINE PRESENT

FIGURE 3 THIN SHADOW LINE DETECTOR  
(Shadow Detection)



FIGURE 4 ORIGINAL IMAGE  
(Shadow Detection)




FIGURE 5 RESULTS OF APPLYING THE LINE  
DETECTOR TO ORIGINAL IMAGE  
(Shadow Detection)




FIGURE 6 NOISE REMOVAL USING MOVING  
WINDOW INTEGRATION ALONG  
SHADOW DIRECTION  
(Shadow Detection)




FIGURE 7 LINE THINNING  
(Shadow Detection)



FIGURE 8 THRESHOLDED LINES OVERLAYED  
ON ORIGINAL IMAGE: COMPARE  
WITH FIGURE 4  
(Shadow Detection)




FIGURE 9 RESULTS USING TWO ADDITIONAL  
IMAGES: COMPARE WITH FIGURE  
10  
(Shadow Detection)




FIGURE 10 RESULTS FOR DETECTING  
EXTENDED OBJECTS  
(Shadow Detection)




FIGURE 11 URBAN SCENE WITH OVERCAST SKY  
(Skyline Delineation)



FIGURE 12 SKY SEED FOUND WITH REGION SEGMENTATION AT THRESHOLD 2  
(Skyline Delineation)




FIGURE 13 THRESHOLD IS DOUBLED UNTIL REGION CONTAINING SKY SEED TOUCHES "BOTTOM" 15% OF PICTURE  
(Skyline Delineation)





FIGURE 14 PICTURE SEGMENTED INTO SKY SEED, UNCLASSIFIED PORTION, AND LAND SEED  
(Skyline Delineation)




(a) ORIGINAL IMAGE




(b) SKY AND LAND SEED BOUNDARIES COINCIDE AT SKYLINE TO  
PRODUCE UNAMBIGUOUS DELINEATION

FIGURE 15 SKYLINE DELINEATION




(a)




(b)

A STEREO PAIR



(c)




(d)


LINE SEGMENTS

FIGURE 16 STEREO MATCHING USING GLOBAL PERSPECTIVE CONSTRAINTS






(e)




(f)

**GAUSSIAN MAPPING OF LINES**  
 (Bright spots indicate vanishing points.  
 The images are also mapped onto the  
 sphere for reference.)




(g)



(h)


**PARALLEL LINES**

**FIGURE 16 STEREO MATCHING USING GLOBAL PERSPECTIVE CONSTRAINTS (Continued)**



(i)

MATCHED LINES



(j)

FIGURE 16 STEREO MATCHING USING GLOBAL PERSPECTIVE CONSTRAINTS (Continued)